Theoretical Computer Science 412 (2011) 5298–5314

Contents lists available at ScienceDirect

Theoretical Computer Science

journal homepage: www.elsevier.com/locate/tcs

Exact and approximate equilibria for optimal group network formation

Elliot Anshelevich [a,1], Bugra Caskurlu [b,*]

[a] Rensselaer Polytechnic Institute, Troy, NY, United States
[b] West Virginia University, Morgantown, WV, United States

ARTICLE INFO

Article history:
Received 29 January 2011
Accepted 30 May 2011
Communicated by P. Spirakis

Keywords:
Algorithmic game theory
Network formation games
Price of stability
Nash equilibrium
Group network formation

ABSTRACT

We consider a process called the Group Network Formation Game, which represents the scenario when strategic agents are building a network together. In our game, agents can have extremely varied connectivity requirements, and attempt to satisfy those requirements by purchasing links in the network. We show a variety of results about equilibrium properties in such games, including the fact that the price of stability is 1 when all nodes in the network are owned by players, and that doubling the number of players creates an equilibrium as good as the optimum centralized solution. For the general case, we show the existence of a 2-approximate Nash equilibrium that is as good as the centralized optimum solution, as well as how to compute good approximate equilibria in polynomial time. Our results essentially imply that for a variety of connectivity requirements, giving agents more freedom can paradoxically result in more efficient outcomes.

1. Introduction

Many modern computer networks, including the Internet itself, are constructed and maintained by self-interested agents. This makes network design a fundamental problem for which it is important to understand the effects of strategic behavior. Modeling and understanding of the evolution of nonphysical networks created by many heterogeneous agents (like social networks, viral networks, etc.) as well as physical networks (like computer networks, transportation networks, etc.) has been studied extensively in the last several years. In networks constructed by several self-interested agents, the global performance of the system may not be as good as in the case where a central authority can simply dictate a solution; rather, we need to understand the quality of solutions that are consistent with self-interested behavior. Much research in the theoretical computer science community has focused on this performance gap and specifically on the notions of the *price of anarchy* and the *price of stability* — the ratios between the costs of the worst and best Nash equilibrium,[2] respectively, and that of the globally optimal solution.

In this paper, we study a network design game that we call the *Group Network Formation Game*, which captures the essence of strategic agents building a network together. In this game, players correspond to nodes of a graph (although not all nodes need to correspond to players), and the players can have extremely varied connectivity requirements. For example, there might be several different "types" of nodes in the graph, and a player desires to connect to at least one of every type (so that this player's connected component forms a Group Steiner Tree [13]). Or instead, a player might want to connect to at least k other player nodes. The first example above is useful for many applications where a set of players attempt to form

* Corresponding author. Tel.: +1 518 2768988; fax: +1 518 2764033.
E-mail addresses: eanshel@cs.rpi.edu (E. Anshelevich), caskurlu@gmail.com (B. Caskurlu).
1 Tel.: +1 518 2766491; fax: +1 518 2764033.

2 Recall that a (pure-strategy) Nash equilibrium is a solution where no single player can switch her strategy and become better off, given that the other players keep their strategies fixed.

0304-3975/$ – see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.tcs.2011.05.049

groups with "complementary" qualities. The second example corresponds to a network of servers where each server want to be connected to at least k other servers so that it can have a backup of its data; or in the context of IP networks, a set of ISPs that want to increase the reliability of the Internet connection for their customers, and so decide to form multi-homing connections through k other ISPs [23]. Many other types of connectivity requirements fit into our framework, and so the results we give in this paper will be relevant to many different types of network problems.

Game definition. We now formally define the *Group Network Formation Game* as follows. Let an undirected graph $G = (V, E)$ be given, with each edge e having a nonnegative cost $c(e)$. This graph represents the possible edges that can be built. Each player i corresponds to a single node in this graph (that we call a *player* or *terminal* node), which we will also denote by i. Similarly to [4], a strategy of player i is a payment vector p_i of size $|E|$, where $p_i(e)$ is how much player i is offering to contribute to the cost of edge e. We say that an edge e is *bought*, i.e., it is included in the network, if the sum of payments of all the players for e is at least as much as the cost of e ($\sum_i p_i(e) \geq c(e)$). Let G_p denote the subgraph of bought edges corresponding to the strategy vector $p = (p_1, \ldots, p_N)$. G_p is the outcome of this game, since it is the network which is purchased by the players.

To define the utilities/costs of the players, we must consider their connectivity requirements. *Group Network Formation Game* considers the class of problems where the players' connectivity requirements can be compactly represented with a function $F : 2^U \to \{0, 1\}$, where $U \subseteq V$ is the set of player nodes, similar to [14]. This function F has the following meaning. If S is a set of terminals, then $F(S) = 1$ if and only if the connectivity requirements of all players in S would be satisfied if S is the set of terminals of a connected component in G_p. For the example above, where each player wants to connect to at least one player from each "type", the function $F(S)$ would evaluate to 1 exactly when S contains at least one player of each type. Similarly, for the "data backup" example above, the function $F(S)$ would evaluate to 1 exactly when S contains at least $k + 1$ players. In general, we will assume that the connectivity requirements of the players are represented by a monotone "happiness" function F. The monotonicity means that $F(B) \geq F(A)$ when B contains A. This implies that if the connectivity requirement of a player is satisfied in a graph G_p, then it is still satisfied when this player is connected to a superset of the nodes she is connected to in G_p. We will call a set of player nodes S a "happy" group if $F(S) = 1$. While not all connectivity requirements can be represented as such a function, it is a reasonably general class that includes the examples given above. Therefore an instance of our game consists of a graph $G = (V, E)$, player nodes $U \subseteq V$, and a function F that states the connectivity requirements of the players. We will say that player i's connectivity requirement is *satisfied* in G_p if and only if $F(S_i(G_p)) = 1$ for $S_i(G_p)$ being the terminals of i's connected component in G_p. While required to connect to a set of terminal nodes satisfying her connectivity requirement, each player also tries to minimize her total payments, $\sum_{e \in E} p_i(e)$ (which we will denote by $|p_i|$). We conclude the definition of our game by defining the cost function for each player i as:

- $cost(i) = \infty$ if $F(S_i(G_p)) = 0$
- $cost(i) = |p_i|$ otherwise.

In our game, all players want to be a part of a happy group which can correspond to many connectivity requirements, some of which are mentioned above. The socially optimal solution (which we denote by OPT) for this game is the cheapest possible network where every connected component is a happy group, since this is the solution maximizing social welfare.[3] For our first example above, OPT corresponds to the cheapest forest where every component is a Group Steiner Tree, for the second to the Terminal Backup problem [5], and in general it can correspond to a variety of constrained forest problems [14]. Our goals include understanding the quality of exact and approximate Nash equilibria by comparing them to OPT, and thereby understanding the efficiency gap that results because of the players' self-interest. By studying the price of stability, we also seek to reduce this gap, as the best Nash equilibrium can be thought of as the best outcome possible if we were able to suggest a solution to all the players simultaneously.

In the *Group Network Formation Game*, we do not assume the existence of a central authority that designs and maintains the network, and decides on appropriate cost-shares for each player. Instead we use a cost-sharing scheme which is sometimes referred to as "arbitrary cost sharing" [4,11] that permits the players to specify the actual amount of payment for each edge. This cost-sharing mechanism is necessary in scenarios where very little control over the players is available, and gives more freedom to players in specifying their strategies, i.e., has a much larger strategy space. The main advantage of such a model is that the players have more freedom in their choices, and less control is required over them. A disadvantage of such a system, however, is that it does not guarantee the existence of Nash equilibria (unlike more constrained systems such as fair sharing [3]). Studying the existence of Nash equilibria under arbitrary cost sharing has been an interesting research problem and researchers have proven existence for many important games [2,4,11,16,17]. Interestingly, in many of these problems it has been shown that the equilibrium is indeed cheap, i.e., costs as much as the socially optimal network. As we show in this paper, this tells us that in the network design contexts we consider, *arbitrary sharing produces more efficient outcomes while giving the players more freedom.*

Related work. Over the last few years, there have been several new papers using arbitrary cost-sharing, e.g., [2,11,16,18]. Recently, Hoefer [17] proved some interesting results for a generalization of the game in [4], and considered arbitrary sharing in variants of Facility Location.

3 The solution that maximizes the social welfare is the one that minimizes the total cost of all the players.